

Abstract

While scalability in database processing may be achieved gradually and at low cost, to achieve high scalability of several thousand times. This database accelerator holds on an accelerator system copies of location tables, which correspond to primary-key indices, and alternate-key location tables, which correspond to alternate-key indices, for a primary system that performs data updates. When a location table or alternate-key location table on the primary system is modified, updates are synchronized to the location table and alternate-key location tables on the accelerator system. The data itself is held only on the primary system. Major improvements in processing capabilities are achieved over serial processing by distributing processing requests among accelerator systems.